

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 747 808 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
11.12.1996 Bulletin 1996/50

(51) Int Cl.⁶: G06F 9/318

(21) Application number: 96480073.4

(22) Date of filing: 31.05.1996

(84) Designated Contracting States:
DE FR GB

(30) Priority: 07.06.1995 US 480105

(71) Applicant: International Business Machines
Corporation
Armonk, N.Y. 10504 (US)

(72) Inventors:
• Goetz, John W.
Jericho, Vermont 05465 (US)

• Mahin, Stephen W.
Underhill, Vermont 05489 (US)
• Bergkvist, John J.
Williston, Vermont 05495 (US)

(74) Representative: de Pena, Alain
Compagnie IBM France
Département de Propriété Intellectuelle
06610 La Gaude (FR)

(54) Processor capable of supporting two distinct instruction set architectures

(57) A microprocessor which supports two distinct instruction-set architectures. The microprocessor includes a mode control unit which enables extensions and/or limitations to each of the two architectures and controls the architectural context under which the microprocessor operates. The control unit controls memory management unit (MMU) hardware that is designed to allow address translation to take place under the control of a mode bit so that the translation mechanism can be switched from one architecture to another. A single MMU translates addresses of the two distinct architectures under control of the mode bit which is also used to simultaneously inform instruction decode which architecture is being used so that instructions are properly decoded. The MMU is also capable of mapping the address translation of one architecture onto that of the other so that software written for both architectures may be multi-tasked under the control of a single operating system.




FIG.8

EP 0 747 808 A2

Description

BACKGROUND OF THE INVENTION

Field of the Invention

The present invention generally relates to microprocessors and, more particularly, to a microprocessor with an architecture mode control capable of supporting extensions to two distinct instruction-set architectures (hereinafter referred to simply as "architectures").

Background Description

There are currently two competing microprocessor architectures used in personal computers. One, referred to as the X86 architecture, was developed by Intel Corporation and supports a family of microprocessors, especially the 80386, 80486, Pentium®, and P6™ processors. The other, referred to as the PowerPC™ architecture, was jointly developed by International Business Machines Corporation and Motorola and currently includes a plurality of PowerPC processors. (PowerPC is a trademark of the IBM Corp.) The PowerPC processors are reduced instruction set computer (RISC) processors, while the X86 architecture is an example of a complex instruction set computer (CISC) architecture.

There is a need to be able to support a broad range of software including a very large installed base of software for the X86 architecture and newer software written to take advantages of the PowerPC processing power. A processor that implements the two architectures does so to allow execution of software from either architecture, thus expanding its potential marketplace. The usefulness and potential marketability of such a processor is greatly enhanced if it can allow either software standard to run dynamically in a multitasking environment. There are, however, several problems with allowing the two architectures to operate in this fashion. These include the different instruction sets and the different use of address space inherent to the two architectures. One approach is to perform a software emulation of the CISC architecture on the RISC processor; however, this approach sacrifices the processing speed of the RISC processor and does not support multitasking software written for the two different architectures. What is needed is a hardware solution that will realize the full potential of the RISC processor speed yet insure full compatibility between the two architectures.

A RISC processor, such as the PowerPC, has a very limited instruction set, whereas modern CISC processors have a very extensive instruction set, including both simple and complex instructions. Adding the instruction set of a CISC processor to a RISC processor would defeat the very purpose of the design of the RISC processor. Therefore, translating the CISC instruction set so that the RISC processor can execute the CISC instructions with its limited instruction set is a formidable prob-

lem. This problem is further exacerbated by the need to distinguish between instructions written for the two different instruction sets so that proper decoding can be accomplished. If multitasking for programs written for the two different architectures is to be accomplished, properly identifying and decoding instructions written for the two different instruction sets must be done dynamically and transparently to the user.

Supporting two architectures on a single processor and allowing dynamic multitasking between software implemented for either architecture also involves the control over the architectural context in which the processor's execution units operate. The X86 and PowerPC architectures, for example, differ greatly in the design not only of the instruction sets but also on the assumptions each instruction places on execution resources, like registers and result flags. These resources also consume considerable space in a processor. Operand size and type, allowable operations, and synchronization requirements of operations also differ between each architecture.

The PowerPC architecture defines a set of thirty-two general purpose registers (GPRs) used in fixed-point operations and a separate set of thirty-two floating-point registers (FPRs) used in floating-point operations. Any of the thirty-two registers may be used in the respective fixed or floating-point operations. Values held in the registers are always right-justified. Specific instructions are defined to load and store data between the registers and memory, and a separate set of instructions are defined to operate on data in the registers. No instructions are defined to, for example, load data from memory and add it to a value in a register. Two separate instructions would be required to perform the operation.

The X86 architecture defines a set of eight GPRs and eight FPRs. The FPRs are organized as a stack rather than a register file. Certain instructions place restrictions on how the registers in the GPRs may be used, making the GPRs less than general. For example, move string operations restrict the use of EDI and ESI as index registers. In four of the GPRs, values are not required to be right-justified; rather, they may be referenced directly from the second byte of the register. Many instructions may perform operations on memory locations. For example, the add instruction may take one source from a GPR, another from memory and write the result back into the source memory location.

PowerPC fixed and floating-point execution instructions often define three register operands: two source operands and a target. Similar X86 instructions define just two operands: a source operand and a source/target operand. One or both of the X86 operands may be memory locations and not just registers, unlike the operands included in instructions for PowerPC processors.

A number of other differences exist between the two architectures' execution resource assumptions that further distinguish them. For example, the PowerPC archi-

itecture defines eight result control fields within a single register, and compare and branch instructions may operate on any of the eight fields, allowing optimizing compilers great flexibility in code generation. The X86 architecture defines only one set of result controls for use by comparison and branch (jump) instructions.

A further problem directly related to the differing instruction set architectures is control over the architectural context in which instructions themselves are fetched from main memory and decoded by the processor's fetch and decode logic. In this case again, the X86 and PowerPC architectures differ greatly in the requirements they place on such logic.

The PowerPC architecture defines all instructions to be exactly four bytes long, with opcodes and operand information in fixed locations. Instructions are always aligned on a word (4-byte) boundary, and as a result never cross cache, page or segment boundaries.

The X86 architecture, on the other hand, defines variable-length instructions, with operand information in non-regular locations dependent on the instruction and any opcode prefixing. Instructions are aligned only at byte boundaries and therefore may cross cache, page and segment boundaries. As a result, the demands placed on the design of instruction fetch and decode logic for each architecture are vastly different, with X86 demands being much more difficult to implement efficiently than PowerPC demands.

Another problem is memory management since the underlying operating system is capable of managing only one of the two architectures' address translation mechanisms. There are two significant reasons for this. One is that the virtual memory management portion of the operating system (referred to as the VMM) is written for just one of the two architectures. The other is that existing processors contain only one memory management unit (MMU) in direct support of the processor's instruction set architecture. MMUs tend to consume significant physical space on a processor die, so physical space constraints impose an additional impediment to implementing a single processor which supports two architectures.

Both the PowerPC and X86 processor architectures define memory management schemes whereby a large virtual memory space may be mapped into a smaller physical address space. In both architectures, the translation from virtual to physical address is a two-step process. First, the effective address calculated as part of instruction execution undergoes segment translation to form a virtual address. The virtual address is then translated via the paging mechanism to form a physical address. While this is the basic process, the terminology sometimes varies. For example, the X86 literature sometimes refers to the effective address as the offset portion of a logical address (the selector forms the remaining portion of a logical address) and a virtual address as a linear address. Despite the similarities in the basic address translation process, the details of seg-

ment and page translation between the two architectures differ greatly.

In a 64-bit version of the PowerPC architecture, effective addresses (EAs) are translated to virtual addresses (VAs) via a hashed segment table search. The lower five bits of the effective segment identification (ESID) extracted from the EA are hashed and then concatenated with an address space register to form the real address of the segment table group in memory. The individual segment table group entries are searched until an entry is found whose effective segment ID matches that of the original EA. When found, the virtual segment ID (VSID) is extracted from the segment table group entry and concatenated with the page and byte fields of the original EA to form the VA.

In a 32-bit version of the PowerPC architecture, the upper four bits of the EA are used as an index into one of sixteen segment registers. The VSID is extracted from the segment register and concatenated with the page and byte fields of the original EA to form the VA.

In X86 architecture address translation, EAs are translated to VAs via a direct segment table lookup. A selector value taken from one of six registers is used as a pointer into one of two descriptor tables. The descriptor table entry pointed to by the selector contains a base address which is added to the original EA to form a VA. The X86 reference material usually refers to the EA as an "offset", the combination of the selector and EA as a "logical address", and the VA as a "linear address".

Page translation is also different in the two architectures. In the PowerPC architecture, VAs are translated to physical addresses (PAs) via a hashed page table search. The lower thirty-nine bits of the virtual segment ID plus the page field from the VA are hashed and then masked/merged with a page table origin register to form the real address of the page table group in memory. The individual page table group entries are searched until an entry is found whose virtual segment ID matches that of the original VA. When found, the real page number is extracted from the page table group entry and concatenated with the byte field of the original VA to form the 64-bit PA.

In the X86 architecture, the VAs are translated to PAs via a direct, two-level page table lookup. The high order ten bits of the VA are used as a pointer into a page directory table whose base is determined by a page directory register. The entry in the page directory table referenced by the VA contains the base address for a page table. The middle ten bits of the VA are used as a pointer into this page table. The page table entry referenced by this pointer contains the real page number of the physical page in memory corresponding to the virtual page being translated. The real page number is combined with the offset field of the VA to form the final PA.

Finally, a dual-architecture multitasking processor must be able to manage the context in which external and asynchronous interrupts are taken as well as any synchronous exceptions, or faults. The X86 and Power-

PC architectures define two different schemes for interrupts and exceptions.

The PowerPC architecture defines a single location in which control is transferred as a result of an external interrupt. The burden is placed on software to determine what the vector number of the interrupt is by querying the system. All interrupts and faults are taken in real-address mode in one of two possible memory locations. The return address is stored in a register for use by the interrupt return instruction.

SUMMARY OF THE INVENTION

It is therefore an object of the present invention to provide a processor which supports two distinct architectures under a single multi-tasking operating system.

Another object of the invention is to provide a single microprocessor which enables extensions and/or limitations (restrictions) to each of two distinct architectures and controls the architectural context under which the processor operates.

It is another object of the invention to provide a microprocessor capable of qualifying the instruction set definitions for two supported architectures so that resources defined in each architecture may be accessed by software written for the other architecture.

It is a further object of the invention to provide a microprocessor which determines an architectural context under which execution resources operate such that the execution context may be dynamically switched from one architecture to another.

It is yet another object of the invention to provide a microprocessor that supports two distinct architectures and which has a mode control unit that initializes the microprocessor in a known state from which software may access various mechanisms to enable/disable a qualifying mode control and influence an architectural context control mechanism.

It is still another object of the invention to provide a microprocessor that supports two distinct architectures and which has memory management hardware capable of performing address translation from virtual to real addresses for both architectures and is designed to allow address translation to take place such that the translation mechanism may be switched from one architecture to another.

It is still another object of the invention to provide a microprocessor that supports two distinct architectures and which has memory protection checking hardware capable of performing memory protection checks for both architectures and allows memory protection checks to take place in a manner such that memory resources of one architecture may be protected from memory resources of the other architecture.

It is yet a further object of the invention to provide a microprocessor that supports two distinct architectures and which determines the architectural context under which interrupts and exceptions are taken such that

the interrupt context may be switched from one architecture to another.

According to the invention there is provided a microprocessor which runs under a single multitasking operating system and supports first and second architectures having separate and distinct instruction sets and memory management schemes. The microprocessor comprises instruction set management means that decodes instructions in a first instruction set of the first architecture and decodes instructions of a second instruction set of the second architecture. The instruction set management means maps decoded instructions in the first instruction set to one or more instructions in the second instruction set. The microprocessor further includes memory management means that performs address translation from virtual to real addresses for said first and second architectures. Control means detects an architectural context of a program being read from memory as being either code for said first architecture or code for said second architecture and, depending on the detected architectural context, controls the instruction set management means and said memory management means to dynamically switch between address translation for the first or second architectures and executing one or more mapped decoded instructions or directly decoded instructions of the second architecture.

In a specific implementation of the invention, the microprocessor is provided with an architecture mode control unit which enables extensions and limitations to the two architectures. These extensions and limitations allow the following:

- Enablement of new instructions and extensions to existing instructions in one architecture to allow full access into unique resources of the other architecture.
- Full visibility of one architecture into the resources of another architecture.
- A single address translation mechanism in effect such that the translation of addresses for one architecture may be mapped onto the translation of another architecture.
- A mapping of the protection mechanism of one architecture onto that of the other architecture.
- A unified interrupt and exception mechanism that allows asynchronous interrupts and page translation and protection related exceptions to be handled by a single mechanism regardless of the architectural context in effect when the interrupt or exception occurred.

Additionally, the architecture mode control unit controls the architectural context (context control) under which the processor operates by controlling the follow-

ing areas of the processor;

- There is a single instruction fetch mechanism shared by the two architecture modes, and separate instruction decode mechanisms for each architecture which are active only when appropriate for the given context control. However, such a context control can be used on implementations with multiple instruction fetch mechanisms and/or single, multi-architecture decoders.
- All execution resources are common between the two architecture modes. However, such a context control can be used on implementations without shared or common resources. For example, an implementation may have separate X86 and PowerPC architecture register files. The context control would be used to select the appropriate register file for operand and result accesses.
- A single memory management unit (MMU) is implemented using a format common to the two supported architectures. However, such a context control can be used on implementations with multiple MMUs to drive translations through the MMU appropriate to the architecture given by the context control. The page protection mechanism to be used by the processor MMU when protecting supervisor-level code from user-level code.

BRIEF DESCRIPTION OF THE DRAWINGS

The foregoing and other objects, aspects and advantages will be better understood from the following detailed description of a preferred embodiment of the invention with reference to the drawings, in which:

Figure 1 is a block diagram of a microprocessor on which the invention may be implemented;

Figure 2 is a block diagram showing the effective to virtual address translation (segmentation) in the PowerPC architecture;

Figure 3 is a block diagram showing the effective to virtual address translation (segmentation) in the X86 architecture;

Figure 4 is a block diagram showing the virtual to physical address translation (paging) in the PowerPC architecture;

Figure 5 is a block diagram showing the virtual to physical address translation (paging) in the X86 architecture;

Figure 6 is a logic diagram showing the relevant logic of an exemplary mode control unit in Figure 1;

Figure 7 is a state diagram showing how the architecture mode control unit generates the architecture context control;

Figure 8 is a high level block diagram showing the data flow of the microprocessor's instruction unit;

Figure 9 is a block diagram of the microprocessor's instruction unit;

Figure 10 is a block diagram contrasting the translation lookaside buffer (TLB) formats for the PowerPC and X86 page table entries;

Figure 11 is a block diagram showing X86 address translation is mapped to PowerPC address translation according to the invention;

Figure 12 is a set of tables and block diagrams showing the PowerPC paged memory protection checking rules;

Figure 13 is a set of tables and block diagrams showing the X86 paged memory protection checking rules;

Figure 14 is a block diagram showing the PowerPC interrupt status and control registers and interrupt vector table;

Figure 15 is a block diagram showing the X86 real mode and protected mode interrupt vector tables; and

Figure 16 is a block diagram showing the X86 interrupt status and control registers and interrupt stack.

DETAILED DESCRIPTION OF A PREFERRED EMBODIMENT OF THE INVENTION

The invention will be described, by way of example, as a specific implementation of a RISC processor, the IBM PowerPC family of processors, modified to support the memory management scheme and instruction set of a CISC processor, the Intel X86 family of processors. It will be understood, however, that the invention could be applied to other and different processors. Moreover, the teachings of the invention may be applied to a combination of a pair of RISC processor architectures or a combination of CISC processor architectures, each having different memory management schemes and instruction sets. Those skilled in the art will also recognize that the invention can be extended to support for multiple processor architectures.

Referring now to the drawings, and more particularly to Figure 1, there is shown a block diagram of the basic microprocessor, such as the PowerPC microprocessor, on which the present invention may be implement-

ed. The following discussion provides the fundamental operation of the microprocessor.

The microprocessor 10 is connected via its system interface 101 to a system bus 12 comprising a 64-bit data bus 121 and a 32-bit address bus 122. The system bus 12 is connected to a variety of input/output (I/O) adapters and a system memory (not shown). The microprocessor 10 uses the system bus 12 for performing reads and writes to system memory, among other things. Arbitration for both address and data bus mastership is performed by a central, external arbiter (not shown).

The system interface 101 is connected to a memory unit 102, which consists of a two-element read queue 1021 and a three-element write queue 1022. The read queue 1021 contains addresses for read operations, and the write queue 1022 contains addresses and data for write operations. The memory unit 102 is, in turn, connected to and receives addresses from a memory management unit (MMU) 103. The memory unit 102 is also connected to a cache 104 which stores both instructions and data. Instructions and data (operands) in cache 104 are accessed by the instruction unit 105, consisting of an instruction queue 1051, program counter 1052, issue logic 1053, and branch prediction unit (BPU) 1054 having a branch history table (BHT).

The issue logic 1053 determines the type of instruction and dispatches it to a corresponding one of a plurality of execution units, here represented by an integer unit (IU) 106 and a floating point unit (FPU) 107. The IU 106 includes an arithmetic logic unit (ALU) 1061 which performs scalar (i.e., integer) operations and stores results in a general purpose register (GPR) file 1062. Similarly, the FPU 107 includes an ALU 1071 which performs floating point operations and stores results in a floating point register (FPR) file 1072. The data outputs from each of the GPR file 1062 and the FPR file 1072 are written to cache 104 from where the data is transferred to the memory unit 102 for writing to system memory. In addition to data calculations, the IU 106 also calculates addresses for accessing by the instruction unit 105 and temporarily stores these addresses in a register 1063. The addresses in register 1063, along with addresses output by the BPU 1054 are supplied to the MMU 103.

The instruction unit 105 also processes interrupts (asynchronous events initiated by hardware external to the processor) and exceptions and faults (synchronous events occurring as a result of fetching, decoding or executing an instruction). Interrupts are sent to the microprocessor 10 via the system interface 101 and forwarded to the issue logic 1053. Exceptions and faults (hereafter referred to simply as "exceptions") may be detected by either instruction queue 1051, memory management unit 103, IU 106, or FPU 107 and forwarded to issue logic 1053. The issue logic 1053 prioritizes exceptions and signals them to the branch prediction unit 1054 on instruction boundaries. The branch prediction unit

1054 then changes the location from which instruction unit 105 fetches instructions to that of the appropriate interrupt/exception handler.

Instructions and operands are automatically fetched from the system memory via the cache 104 into the instruction unit 105 where they are dispatched to the execution units at a maximum rate of three instructions per clock. Load and store instructions specify the movement of operands to and from the integer and floating-point register files and the memory system. When an instruction or data access is made, the logical address (effective address) is calculated by the instruction unit 105 (for instruction accesses) or integer unit 106 (for data accesses). The memory management unit 103 translates the effective address to a physical address and forwards that to the cache 104. While translating the effective address to a physical address, the memory management unit 103 also checks the current privilege of the microprocessor 10 to verify that the memory may be accessed. A portion of the physical address bits are compared with the cache tag bits IN 104 to determine if a cache hit occurred. If the access misses in the cache 104, the physical address is used to access system memory.

In addition to loads, stores and instruction fetches, the microprocessor 10 performs other read and write operations for table searches, cache cast-out operations when least-recently used (LRU) sectors are written to memory after a cache miss, and cache-sector snoop push-out operations when a modified sector experiences a snoop hit from another bus master. All read and write operations are handled by the memory unit 102. To maintain coherency, the write queues 1022 are included in snooping. Memory is accessed through an arbitration mechanism that allows devices to compete for bus mastership.

Microprocessor 10 also contains a mode control unit 108 which controls the architectural context under which the various units operate as well as any architectural qualifications or extensions that might be placed on those units under a given architectural context. The memory management unit (MMU) 103 detects to which architecture an instruction conforms. For example, the MMU 103 may determine whether an instruction is an X86 instruction or a PowerPC instruction and informs the mode control unit 108 of the architectural context accordingly via a page mode control (P) bit. The mode control unit 108 also receives a virtual (V) bit from the PowerPC MSR.IR register 109. Based on the P and V bits and a qualification (Q) bit which the mode control unit 108 generates and holds, the mode control unit 108 generates a mode switch signal that determines how instruction unit 105 fetches and decodes instructions and which register resources are available to integer unit 106 and floating point unit 107. Mode control unit 108 also governs how branch prediction unit 1054 redirects instruction unit 105 when interrupts/exceptions occur.

Extensions to an architecture might include new in-

structions and registers. Limitations might include the disabling of access to certain registers, or a restriction on how addresses are translated. The primary purpose behind the extensions and/or limitations is to allow software written in each architecture to have some level of access to resources defined by the other architecture. A secondary purpose is to allow software written in each architecture to transfer control to software written in the other architecture as the architectural context changes. In fact, the extensions enabled by the architecture mode control unit 108 include mechanisms that allow software to invoke a change in architecture context.

The architecture context of the processor is simply the current instruction-set architectural context under which the processor operates. In a processor which supports both the PowerPC and X86 architectures, for example, the context control determines whether the processor behaves like a PowerPC or an X86 processor. Both the architecture qualifying control (which enables extensions/limitations) and the architecture context control (which determines context) may have a direct influence over each of the

- instruction set definition mechanisms,
- instruction encoding mechanism,
- instruction opcode length mechanism,
- address translation mechanism,
- segment and page table organization mechanism,
- protection mechanism,
- interrupt architecture mechanism,
- memory addressability mechanism,
- register sets and register mechanism, and
- conditions, fields and results mechanism.

In the preferred embodiment of the invention, the mode control unit 108 is controlled by a single MMU 103 which is capable of mapping the address translation of one architecture onto that of the other. The mode control unit 108 further controls memory protection checking hardware so that software written for both architectures may be protected and multi-tasked under the control of a single operating system.

Figure 2 is a block diagram of that part of the PowerPC architecture which performs effective to virtual address translation (segmentation). The 64-bit effective address is held in register 21, while address space register (ASR) 22 holds the real address of the segment table. The lower five bits of the effective segment identification (ESID), bits 31 to 35, extracted from the EA

register 11 are hashed by a hash function 23 and then concatenated with bits 0 to 51 of ASR 22 to form the real address in segment table entry register 24, the last byte of which is forced to zero. The segment table entry register 24 addresses the segment table 25 in memory which comprises 4096 bytes. Effective addresses (EAs) are translated to virtual addresses (VAs) via the hashed segment table search. The individual segment table group entries are searched until an entry is found whose effective segment ID matches that of the original EA. When found, the virtual segment ID is extracted from the segment table group entry 26 and concatenated with the page and byte fields of the original EA to form the 80-bit VA in register 27.

In X86 architecture address translation as shown in Figure 3, EAs are translated to VAs via a direct segment table lookup. A selector value taken from one of six registers 31 is used as a pointer into one of two descriptor tables 32. The descriptor table entry pointed to by the selector contains a base address which is read from descriptor table base register 33 and added to the original EA in EA register 34 by adder 35 to form a VA in VA register 36. The X86 reference material usually refers to the EA as an "offset", the combination of the selector and EA as a "logical address", and the VA as a "linear address".

Page translation is also different in the two architectures. In the PowerPC architecture as shown in Figure 4, VAs in register 27 are translated to physical addresses (PAs) via a hashed page table search. More particularly, bits 52 to 67 (page field) of the 80-bit virtual address in register 27 are concatenated with 23 zeros in register 42. Then bits 13 to 51 of the virtual segment ID in register 27 and the content of register 42 are hashed in hash function 43 to generate 39 bits in register 44. Bits 58 to 63 in hash table register 45 are decoded by decoder 46 to generate a 28-bit mask in register 47. Bits 0 to 27 of register 44 are masked by the mask in register 47 in AND gate 48 and the masked output is merged with bits 18 to 45 of register 45 in OR gate 49 to form the mid 28 bits in page table origin register 50. The first 18 bits of register 50 are read directly from bits 0 to 17 of register 45, and the next higher 11 bits are read from bits 28 to 38 of register 44. The highest seven bits of register 50 are forced to zero to form the real address of the page table 51 in memory. The individual page table group entries are searched until an entry 52 is found whose virtual segment ID matches that of the original VA. When found, the real page number is extracted from the page table group entry 52 and concatenated with the byte field of the original VA, bits 68 to 79 of VA register 27, to form the 64-bit PA in physical address register 53.

In the X86 architecture as shown in Figure 5, the VAs are translated to PAs via a direct, two-level page table lookup. More particularly, the high order ten bits of the VA in VA register 36 are used as a pointer into a page directory table 54 whose base is determined by a

page directory register 55. The entry in the page directory table 54 referenced by the VA contains the base address for a page table 56. The middle ten bits of the VA in register 36 are used as a pointer into this page table 56. The page table entry referenced by this pointer contains the real page number of the physical page in memory corresponding to the virtual page being translated. The real page number is combined by adder 57 with the offset field of the VA in register 36 to form the final PA in physical address register 58.

In order to achieve the implementation of two architectures that can run on a single microprocessor, there are several essential elements that need to be controlled by the mode control unit 108 shown in Figure 1. Specifically, the mode control unit 108 handles cases where we have two architectures which differ in

- instruction set definition, encoding, and opcode length,
- address translation,
- segment and page table organization,
- protection mechanism,
- interrupt architecture,
- memory addressability,
- register sets, and
- conditions, fields and results mechanism.

In all areas in which the architectures differ, the processor implementation may contain hardware that individually supports the differing aspects of each architecture, or it may contain common hardware resources which are capable of supporting all architectures, to eliminate redundancy of hardware resources. In situations where distinct hardware elements are used to support the architectures, the architecture mode control unit 108 is responsible for enabling the appropriate hardware element and disabling the other elements. For example, if the processor implements separate instruction set decoders for each of the supported architectures, the architecture mode control unit will enable the instruction set decoders for the architecture currently in use and disable the decoders for the remaining architectures. In situations where common hardware resources are utilized to eliminate redundancy, the architecture mode control unit will direct such hardware to operate under the rules of the current architecture context, when appropriate, making use of any extensions defined for that context. For example, if the processor implements a common register file in order to minimize physical hardware and wiring, then the architecture mode control unit will control which registers in the common register

file may be accessed under any given hardware context.

Therefore, according to the invention, the mode control unit 108 selects between hardware resources which individually support the differing aspects of a single architecture. The control unit 108 also controls those common hardware resources which are capable of supporting the differing elements of multiple architectures, those hardware resources being implemented to eliminate redundant hardware thereby saving physical space, power consumption, improving processor performance, and simplifying the overall design.

The architecture qualifying control mechanism is determined by the value of a bit held in processor feature control register (shown in Figure 6) in the mode control unit 108. This bit may be set by software running in either architecture. When this bit is zero, no qualifications are placed on either of the architectures, e.g., the PowerPC architecture or the X86 architecture; that is, this is the normal state of operation. When this bit is one, however, it qualifies the architecture presently running under the control of the architecture context control mechanism. More particularly, when this bit is one, the microprocessor has available new instructions in the current architecture which allow the microprocessor to read and write registers in the other architecture that are not defined by the current architecture.

As an example, the X86 architecture has a status/control register called FLAGS which is not present in the PowerPC architecture. So a new PowerPC instruction is defined that allows the microprocessor to read and write the FLAGS register, and this instruction may be used only when the qualifying control bit is one. Likewise, this bit disables some instructions in the current architecture. As an example, when the qualifying control signal is a one, the PowerPC architecture no longer recognizes the X86 architecture INVLPG (invalidate page table entry in the TLB) instruction, when the architecture context is X86 mode. When the qualifying control mechanism is one and the context control mechanism is zero (PowerPC architecture mode), extensions to the PowerPC architecture are enabled that give it access to the X86 architected resources.

For example, enhancements are made to the instruction fetch mechanism of the PowerPC architecture that allow software designed for a PowerPC processor to branch to byte-aligned X86 software. Additional PowerPC processor instructions are enabled that allow software designed for a PowerPC processor to read and write the 64-bit X86 descriptor registers. Finally, extensions to the PowerPC processor paging mechanism are enabled that allow an operating system to differentiate memory locations that contain code designed for a PowerPC processor from those locations that contain X86 code.

Figure 6 is a logic diagram showing the relevant logic of an exemplary mode control unit 108. The hardware feature control register 61 is a latch which is set by software. When the latch is set, it enables AND gate

62 which outputs the mode switch signal. The mode switch signal is generated depending on states of the V bit, the P bit and the Q bit. The V (virtual mode) bit is input to the mode control unit 108 via the PowerPC MSR.IR register 109 and indicates whether PowerPC instruction address translation is enabled (V=1) or disabled (V=0). The P (page mode control) bit is supplied by the MMU 103 and indicates whether the current page is a PowerPC page (P=0) or an X86 page (P=1). The Q bit is the architecture qualifying bit. The P bit and the V bit are input to AND gate 63, the output of which is supplied to one input of exclusive OR (XOR) gate 64. The second input of XOR gate 64 is supplied by inverter 65 connected to register 66 which stores the Q bit. The output of register 66 is supplied directly to one input of multiplexer (MUX) 67 and, via inverter 68, to the other input of MUX 67. The MUX 67 is controlled by the mode switch signal from AND gate 62.

The initial values (i.e., at hardware reset) are V=0, P=X, Q=0, and register 61=0, where X means "don't care". The mode switch signal is enabled by setting register 61 to one. When the architecture qualifying bit, Q, is one and the hardware feature control register 61 is also set to one (X86 architecture mode), certain limitations are placed on the X86 address translation mechanism. Specifically, the X86 paging function is disabled, and it is replaced by the full PowerPC address translation mechanism of segment and page translation. The X86 segment translation mechanism is still utilized.

Certain other paging-related operations are also disabled. Specifically, X86 writes to the page directory base register have no effect, the X86 page translation mode cannot be enabled, and X86 software cannot invalidate entries held internally by the MMU 103 in Figure 1. Instead, these operations are trapped internally by the processor and handled as appropriate. Extensions to the X86 architecture are also enabled by the architecture qualification mechanism state bit. Specifically, X86 instructions may access certain PowerPC architected registers.

Architectural context control in the microprocessor is determined by the value of the Q bit generated and by the mode control unit 108 and held in register 66. For example, when this bit is zero, the microprocessor follows the full set of rules associated with the PowerPC architecture, from instruction fetch and decode to execution to address translation to interrupt and exception handling. When this bit is one, the processor follows the full X86 architectural rules. In the preferred embodiment, the initial value of this bit after reset is one, placing the processor in a mode in which the full X86 architecture is followed.

The value of the Q bit may be changed by the architecture mode control unit 108 to place the microprocessor in PowerPC mode while running in X86 mode, or X86 mode while running in PowerPC mode. This is accomplished by MUX 67. Figure 7 shows a state diagram of how the preferred embodiment's mode control unit

108 may change the value of the architectural context control signal. Figure 7 shows four states; full X86 architecture mode, qualified X86 architecture mode, qualified PowerPC architecture mode, and full PowerPC architecture mode. In the preferred embodiment, the mode control unit 108 uses the value of the V, P and Q bits to advance the processor from one architectural context to another. As mentioned, the Q bit is held in register 66 and may be set by software running in either of the two architectures. The bit V is also under software control and is cleared on exception. The bit P is supplied to mode control unit 108 by the MMU 103.

In the preferred embodiment, a reset of the microprocessor from any of the architectural contexts will put the processor in full X86 mode. That is to say, if the microprocessor is reset, the context and qualifying control mechanisms are put to some known and initial state. In the preferred embodiment, a reset leaves the processor in X86 context with no qualifications. If Q (see Figure 7) is never set to one, the processor will never leave full X86 mode. However, if V, P and Q are all set to one, the microprocessor will enter qualified X86 mode, and will stay in that mode as long as V, P, and Q remain one. Qualified PowerPC mode may be entered from either full X86 mode or qualified X86 mode using one of two methods; either by setting V to zero while Q is one, or by setting V and Q to one while setting P to zero. Full PowerPC mode may be entered from either of the qualified modes if Q is set to zero.

Figure 8 is a high level, functional block diagram showing the data flow of the issue logic 1053 (Figure 1) to illustrate the instruction set management operation of the microprocessor. This is the data flow of the hardware shown in more detail in Figure 9, described in more detail hereinafter. On power up or reset, the microprocessor enters an initialization mode, which begins by assuming a default mode. In the preferred embodiment, the default mode is the full X86 mode, as shown in Figure 7, which is hardwired into the mode control unit 108. The MMU 103 is initially clear. On power up or reset, instruction code is retrieved via the system interface 101 and bus 12 from main memory. This code is supplied to the instruction unit 105 via the normal data path described above.

The instruction code is supplied in parallel to decode function block 70, which handles simple PowerPC instructions, decode function block 71, which handles simple X86 instructions, decode function block 72, which handles complex PowerPC instructions, and decode function block 73, which handles complex X86 instructions, as shown in Figure 8. In the example being described, "simple" instructions are those which map to a basic operation class and can be handled by a single execution unit. For example, a load operation, a store operation and single arithmetic operations are all simple instructions. All other instructions are, by this definition, "complex". For example, an X86 repeat move string or a PowerPC load multiple word are examples of complex

instructions. The outputs of decoding function 70, i.e., decoded simple PowerPC instructions, are supplied to one input of multiplexer (MUX) 74. The outputs of decoding function 71, i.e., decoded simple X86 instructions, are first supplied to a translation unit 75 (e.g., a table lookup ROM) which outputs corresponding simple PowerPC instructions to the second input of MUX 74. The outputs of decoding functions 72 and 73 are supplied to a multiplexer (MUX) 76, the output of which is supplied to a microcode read only memory (ROM) 77.

The ROM 77 has tables of instructions for both complex decoded PowerPC and X86 instructions. For complex PowerPC decoded instructions, the ROM 77 outputs multiple, simple PowerPC decoded instructions. A complex decoded X86 instruction output from MUX 76 is mapped to corresponding multiple, simple PowerPC instructions in ROM 77. Thus, the outputs of MUX 74 and the ROM 77 are one or more simple PowerPC instructions. The context from mode control unit 108 (Figure 1) controls the multiplexers 74 and 76 to select either PowerPC or X86 decoded instructions output, respectively, by the decoding functions 70, 72 and 71, 73. The outputs of MUX 74 and the ROM 77 are supplied to a third MUX 78, which is controlled by OR gate 79. The OR gate 79 receives a "valid" output from one of the decoding functions 72 or 73 when a "complex" instruction is detected by that function. Thus, the output of the OR gate selects between complex and simple instructions; that is, between the output of ROM 77 and the MUX 74.

In order to set up the X86 and PowerPC qualified modes (see Figure 7), the initialization software performs the following steps:

- set the hardware feature control register to one in order to perform an immediate context switch to Power PC mode from the default X86 mode,
- initialize the PowerPC virtual environment (i.e., set up the necessary page tables in external memory and initialize the BATs and segment registers),
- enable instruction relocation by setting the IR bit in the MSR register to one, thus enabling virtual translation of instruction addresses, and
- performing a branch instruction to the appropriate software mechanism that will handle application start up (e.g., either X86 or PowerPC code).

After hardware reset, the processor will be executing in X86 mode (the default); therefore, the Q bit will be set to zero by the mode control unit 108, and the hardware feature control register will also be initialized to zero. Initialization software should subsequently set the hardware feature control register to one in order to enable context switching. Since instruction relocation is off (i.e., the V bit is zero), this will also force the mode control unit 108 to perform a context switch from executing

X86 code to executing PowerPC code. The next sequential instruction that follows the feature control register write will be executed as a PowerPC instruction.

Next, during the initialization process, the page table 52 (Figure 4) is set up and entries made to enable the qualified PowerPC mode (Figure 7) that will become effective once instruction relocation is enabled (i.e., the MSR IR bit is set to one, resulting in the V bit being set to one). This also requires the PowerPC segment registers and BAT registers to be initialized to their appropriate values. Then, the X86 descriptor tables should be initialized for use by X86 applications. Finally, instruction relocation would be enabled by setting the IR bit in the PowerPC MSR register to one, thus setting the V bit to one. The last instruction would be a branch instruction to a memory location that starts the first program. The target code can be either X86 or PowerPC code. Which ever it is, the code is managed in 4KB blocks.

When the branch instruction executes, the values for Q and V will be Q=1 (executing PowerPC code) and V=1 (instruction relocation/virtual translation enabled). Hence, the subsequent state of the Q bit will be determined by the mode control unit 108 from the accessed value of the P bit from the page table entry corresponding to the target of the branch instruction, the value of the P bit being provided to the mode control unit 108 from the MMU 103. At this point, the software initialization is complete.

Once initialization is complete, the MMU 103 performs a TLB lookup. If the lookup results in a hit and the P bit is a one, the MMU 103 knows that the code is X86 code and informs the mode control unit 108 and returns the status via the page mode control bit, P, thus directing the mode control unit as to the context mode. What can be seen from this flow is that the incoming instruction causes the MMU 103 to switch the mode of the mode control unit 108 when the instruction is located in the MMU 103 or added to the MMU. In other words, the addressing scheme of the MMU 103 drives the mode control unit 108 to change in mode. If, however, the incoming instruction is not in the TLB (a miss), a page table walk is performed to find the instruction in main memory lookup table. If it is not there either, a special interrupt is made and handled.

Assuming that the instruction is in the page table, it is loaded in the MMU 103 and, if the P bit is different from the current context mode, the MMU 103 switches the mode of the mode control unit 108, as described above. The MMU 103 always overrides the mode control unit context mode. If this instruction is the next instruction to be executed, then the mode switch occurs. The switch between the qualified PowerPC mode to the qualified X86 mode (Figure 7) is made by informing the mode control unit 108 via the P bit (P=1). Now the mode control unit 108 tells the integer unit (IU) 106 and the floating point unit (FPU) 107 that all instructions are to be interpreted as X86 instructions. More particularly, the address translation mechanisms and the segment and

page table organization mechanisms are in the IU 106. The mode control unit 108 tells the IU 106 how many registers there are and what format the registers use. The mode control unit 108 tells the FPU 107 whether to use 80-bit X86 style or 64-bit PowerPC style, thereby controlling the precision of the same register files.

Figure 8 shows how the instruction stream comes into the instruction unit and how it is decoded. The mode control unit controls the selection of decoded instructions through the issue logic. Figure 9 is a detailed hardware block diagram of the instruction unit 105 which performs the functions illustrated in the functional block diagram of Figure 8. The instruction queue 1051 in Figure 1 includes the fetch aligner 801 and the instruction queue 802 shown in Figure 9. Fetch aligner 801 is connected to the instruction fetch bus and has as inputs instructions and the current instruction address taken from the instruction cache. The architecture mode control signal from mode control unit 108 is used to determine how instructions are aligned by the fetch alignment hardware 801 prior to being transferred to the instruction queue 802 since fetch alignment requirements differ between the PowerPC and X86 architectures. For example, X86 instructions may contain anywhere from one to fifteen bytes and be aligned on any byte boundary whereas PowerPC instructions are always four bytes in length and are therefore always aligned on a 4-byte boundary. The fetch alignment hardware 801 must be able to shift instructions from the instruction fetch bus by the appropriate amount based on the current instruction's length and address. In addition, since X86 instructions may vary in length and may be aligned anywhere on the instruction fetch bus, an X86 instruction may overflow one instruction cache entry and fall into the next entry. The architecture mode control signal is used to allow the fetch alignment hardware 801 to merge and align instructions from consecutive instruction cache entries into the instruction queue 802.

Instructions are forwarded from instruction queue 802 to the X86 instruction decoders 803 and PowerPC instruction decoders 804. The X86 instruction decoders 803 perform the decode functions 71 and 73 of Figure 8, while the PowerPC decoders 804 perform the decode functions 70 and 72. The architecture context control is used to enable either the X86 instruction decoders 803 or PowerPC instruction decoders 804, depending on the desired architectural context, as explained with reference to Figure 8. The decoders are used to translate all instructions into a common set of instructions whose execution is supported by the microprocessor. For example, an X86 "ADD EAX,EBX" instruction and a PowerPC "add r8,r8,r11" instruction will both be translated by their respective decoders to a common "add r8,r8,r11" instruction. More complex instructions will be handled by the respective decoders 803 or 804 which will generate sequences of simple instructions which perform the function of the single, more complex instructions. For example, the PowerPC instruction decoders 804 will

emulate a PowerPC load-multiple word instruction as a sequence of individual loads.

The architecture mode control signal from mode control unit 108 also has a direct effect on the decoders 803 and 804, as described with reference to Figure 9. Specifically, new instructions are enabled in the PowerPC architecture which allow PowerPC software to manage the X86 floating-point environment and to read and write the X86 descriptor registers. In the X86 architecture, the architecture mode control signal disables the X86 instructions which enable page translation, initialize the page directory base register, and invalidate entries in the TLB. The qualifying control signal also enables extensions to existing PowerPC instructions which allow PowerPC software to branch directly to X86 software located anywhere in memory. Specifically, the instructions which branch to addresses stored in the PowerPC branch instructions normally ignore the low-order two bits of the register since all PowerPC instructions are located on 4-byte boundaries. X86 instructions, however, are byte-aligned, so PowerPC branches cannot reach instructions which lie on one, two or three byte boundaries. The qualifying control signal solves this problem by enabling the low-two bits of the branch address registers previously mentioned, thereby allowing PowerPC code to branch to byte-aligned addresses.

The common instructions generated by decoders 803 and 804 under the control of the architecture mode control signal have a direct effect on the operation of execution resources. The X86 decoders 803 may generate common instructions that use only a subset of the registers in GPR 1062 (Figure 1) and FPR 1072 (Figure 1) whereas PowerPC decoders 804 may generate common instructions that make use of all the registers. Further, while the architecture mode control signal controls the context of the two architectures, a number of registers used by one architecture may be accessed by the other architecture. Specifically, PowerPC software may read and write all X86 registers, whereas X86 software may read and write the PowerPC MSR, BATs and decrementer register.

From X86 instruction decoders 803 and PowerPC instruction decoders 804, the decoded common forms of the instructions are transferred to issue logic 806 via instruction selection logic 805, which is controlled by the architecture mode control signal. Additionally, instruction length information is forwarded to the next instruction fetch address (NIFA) compute block 807. The NIFA compute block 807 is used to calculate the address of the next instruction which must be fetched from the instruction cache 104 (Figure 1). It takes as its inputs the address of the current instruction from the instruction fetch bus, the instruction length from the two architecture decoders 803 and 804, and the architecture mode control signal which selects between the two instruction lengths. The branch processing unit (BPU) 808 also has an input into NIFA compute block 807 which allows branch instructions to force the next instruction fetch ad-

dress to the target of the branch.

The architecture qualifying control bit, Q, is input to the memory management unit (MMU) 103 (Figure 1) to control which of the address translation mechanisms is utilized by the microprocessor. The on-chip MMU 103 contains a single page table cache, or translation lookaside buffer (TLB), that is capable of storing either PowerPC or X86 page table entries in a common format. The setting of the architecture mode control signal determines how the TLB entries are initialized by the processor hardware and how they are interpreted during address translation. Figure 10 shows the common format of the microprocessor's TLB entry and how page table entries in the PowerPC and X86 architectures are stored within.

In the present invention, the architecture qualifying control bit, Q, modifies the address translation mechanism used by both architectures. When the architecture qualifying control bit is one, extensions to the PowerPC paging mechanism are enabled that allow an operating system to differentiate memory pages that contain PowerPC code from those pages that contain X86 code. The root of the extension consists of the page mode control bit, P, in each PowerPC page table entry which defines each page as either a PowerPC page (bit set to zero) or an X86 page (bit set to one). This bit is used to drive the value of the architecture context control signal as described above and represented in Figure 7.

As an instruction is fetched from memory, its address is first translated by the MMU 103 shown in Figure 1. The MMU 103 returns status bits with the translated address, including the value of the page mode control bit, P. If the value of the page mode control bit is equivalent to the current value of the architecture qualifying control bit, the processor continues to operate as it did before, in the same architecture. However, if the value of the page mode control bit is different from the architecture qualifying control bit, instruction fetch and decode is halted until all previously fetched instructions are executed and completed without exceptions. Once they have completed, the value of the architecture mode control signal is changed to match that of the page mode control bit, and instruction fetch and decode is restarted under the new context established by the architecture mode control signal.

The enhancements to the PowerPC translation mechanism and the limitations placed on the X86 mechanism specifically allow a mapping of X86 address translation onto PowerPC address translation to provide a more dynamic environment for running software written for both architectures in a multi-tasking operating system. When the architecture qualifying control bit, Q, is set to one, the PowerPC address translation mechanism is "in control". What this means is that X86 segment translation will take place but, instead of performing X86 page translation, PowerPC segment and page translation will follow in order to form a PA. This allows a single operating system to manage address transla-

tion for software written for either architecture. The only restriction is that X86 software can not perform X86 paging-related work of its own.

A 64-bit PowerPC implementation will translate 64-bit addresses, while the X86 architecture as defined today generates 32-bit addresses. The invention provides a means for generating 64-bit addresses from 32-bit X86 addresses by concatenating a 32-bit register value as the high order 32 bits of the final 64-bit address. This allows the 32-bit X86 address space to be located anywhere in 64-bit space and, by dynamically changing the value of the register, allows an operating system to manage multiple X86 address spaces. In 32-bit implementations, this register value is ignored (i.e., it is effectively forced to zero).

Figure 11 shows at a high-level how the two address translation mechanisms are merged using the method according to the invention and where in the process the 32-bit register value is concatenated to the X86 address in order to form a 64-bit address. Register 1001 holds the 32-bit offset and the 16-bit selector which comprise the X86 logical address. The selector is used to address the descriptor table 1002. The segment descriptor from table 1002 is combined in adder 1003 with the offset to generate the 32-bit X86 linear address in register 1004. The 32-bit X86 base address in register 1005 is read into the base 32 bits of the 64-bit PowerPC effective address register 1006. The 32-bit linear address in register 1004 is concatenated with the 32-bit base address in register 1006 to form the 64-bit PowerPC effective address. The effective segment ID from register 1006 is used to address the segment table 1007 to generate the 80-bit PowerPC virtual address in register 1008, as described with reference to Figure 2. The virtual segment ID of register 1008 is used to address the page table 1009 to generate the 52-bit PowerPC real address as describe with reference to Figure 4.

The memory management unit 103 of Figure 1 also is responsible for performing paged memory protection checks in the PowerPC and X86 architectures under the control of the architecture context control signal. Figure 12 shows the PowerPC page protection mechanism. Protection information is gathered from three sources. The Ks and Kp segment protection bits found in the segment table entry 1101 (64-bit PowerPC) or segment register 1101 (32-bit PowerPC), the supervisor/user mode bit (PR) found in register MSR 1102, and the PP page protection bits found in the lower page table entry 1104. A protection key 1105 is formed by ANDing the Kp segment protection bit with the MSR.PR supervisor/user mode bit and ORing that result with the AND of the Ks segment protection bit and the negation of the MSR.PR supervisor/user mode bit. Using the key 1105, the memory management unit 103 (Figure 1) checks the value of the page protection PP bits to determine the type of access allowed, as shown in table 1106.

Figure 13 shows how the X86 page protection mechanism works. Protection information is gathered

from four sources; the descriptor privilege level of the current stack segment descriptor 1201 (this is often referred to as the current privilege level, or CPL), the write-protection bit WP from the register CR0 1202, the user/supervisor bit U/S from the page table entry 1203, and the read/write bit R/W, also from the page table entry 1203. Using this information, the memory management unit performs the check as shown in table 1204 to determine the type of access allowed.

In the present invention, the architecture qualification mechanism bit modifies only the protection mechanism used when the processor is running in X86 mode. Specifically, the X86 page protection mechanism is completely replaced by the PowerPC page protection mechanism since X86 page translation has been replaced by PowerPC page translation. In X86 mode, the qualification mechanism bit causes the PowerPC MSR.PR bit to reflect the privilege of the X86 CPL. An X86 CPL of 0, 1, or 2 forces the MSR.PR to a value of 0 (supervisor) and an X86 CPL of 3 forces the MSR.PR to a value of 1 (user). The protection key 1105 (Figure 11) is then formed as described above. By forcing the MSR.PR to track the CPL in X86 mode, X86 supervisor code may be protected from X86 and PowerPC user code. This allows portions of the operating system to be implemented in X86 code where convenient, since operating system software normally runs at a supervisor level.

Figure 14 is a block diagram of that part of the PowerPC architecture which controls the operation of interrupts and exceptions. Machine state register 1301 is used to enable and disable external interrupts via the EE bit 1310, and it is used to determine the location of the interrupt vector table 1307 in physical memory via the IP bit 1311. Register SRR0 1302 is used to record the effective address of the instruction that resulted in the exception or would have executed in the case of an interrupt. Register SRR1 1302 contains status information about the specific cause of an exception as well as certain bits in the MSR 1301 prior to the interrupt or exception. Register DAR 1304 holds the effective address of the data operand that caused an exception in data-related exceptions. Register DSISR 905 contains status information about the specific cause of a data related exception.

When an interrupt or exception is taken by the PowerPC architecture, the location of the interrupt procedure depends on the type of interrupt and the value of the MSR.IP bit 1311. The MSR.IP bit 1311 specifies whether the interrupt vector table base address 1306 has a value of 0x00000000 or 0xFFFF0000. This address is added to the offset into the vector table 1307 specified by the type interrupt to form the physical address of the interrupt procedure 1309. For example, if IVT base address 1306 is 0xFFFF0000 and a data storage interrupt is taken, the physical address of the interrupt procedure 1309 will be 0xFFFF00300.

Figure 15 shows the two methods in which interrupt

procedure locations may be specified by the X86 architecture. In both methods the interrupt number 1401 is specified by either an instruction, an external interrupt or an internal exception. In X86 real mode (protected mode disabled), the interrupt vector table 1402 usually has a base address of 0x00000000. An interrupt procedure is located by multiplying the interrupt number 1401 by four, thus providing a "far" pointer to the interrupt processor. The "far" pointer is comprised of an offset and a segment value. In X86 protected mode (address translation enabled), the interrupt number 1401 is multiplied by 8 to yield an offset into an interrupt descriptor table 1403. The interrupt gate 1404 referenced by the scaled interrupt number 1401 contains an offset into a destination code segment 1405, and an offset into a segment table 1406. The segment descriptor 1407 in pointed to by the offset into segment table 1406 contains the base address of the destination code segment 1405. This base address is added to the offset into the destination code segment 1405 to form the effective address of the interrupt procedure 1408.

When an interrupt is taken in the X86 architecture, certain information is recorded on the interrupt procedure stack as shown in Figure 16. This consists of the stack pointer of the interrupted procedure (the old SS 1501 and old ESP 1502), a pointer to the instruction that was interrupted (the old CS 1504 and old EIP 1505), an error code 1506, and a copy of the EFLAGS 1503 which contains the state of the interrupted instruction. In addition, page fault interrupts will store the address of the instruction or data byte that caused the page fault in CR2 1507.

Interrupts and exceptions for both the PowerPC and X86 architectures are handled by the branch processing unit (BPU) 1054 in Figure 1. The architecture context mechanism bit from the architecture mode control unit 108 determines which of the two interrupt and exception mechanisms, PowerPC or X86, will be used by the microprocessor.

When the architecture qualification mechanism bit is enabled, interrupts and exceptions encountered by the processor are handled by either the PowerPC or X86 architected mechanism, depending on the type of interrupt or exception as well as the state of the architecture context mechanism bit. Specifically, all asynchronous interrupts are directed to the PowerPC interrupt mechanism regardless of the state of the architecture context mechanism bit. This is done to allow an operating system to have a consistent point of control over events that may have no connection to what the processor was executing when they occurred.

Synchronous exceptions are, in general, directed to the architected mechanism as defined by the architecture context mechanism bit. This is not true, however, in the cases of (a) any form of exception resulting from page translation and (b) any form of exception resulting from PowerPC protection checking taking place within the MMU. In both cases the exception is directed to the

PowerPC mechanism regardless of the value of the architecture context mechanism bit. This is done because X86 paging is not activated, so all PowerPC MMU-related exceptions must be forwarded to the PowerPC mechanism.

While the invention has been described in terms of a single preferred embodiment, those skilled in the art will recognize that the invention can be practiced with modification within the scope of the appended claims.

Claims

1. A processor which supports first and second architectures having separate and distinct instruction sets and memory management schemes, said processor running under a single multitasking operating system and comprising:
 - instruction set management means for decoding instructions in a first instruction set of said first architecture and directly decoding instructions of a second instruction set of said second architecture and for mapping decoded instructions in the first instruction set to one or more instructions in the second instruction set;
 - memory management means for performing address translation from virtual to real addresses for said first and second architectures; and
 - control means for detecting an architectural context of a program being read from memory as being either code for said first architecture or code for said second architecture and controlling said instruction set management means and said memory management means to dynamically switch between address translation for the first or second architectures and executing one or more mapped decoded instructions or directly decoded instructions of the second architecture.
2. The processor recited in claim 1 wherein the memory management means comprises a mode control mechanism controlled by said control means for controlling which of first or second architectural translation methods is to be used by a memory management unit when translating an effective address to a virtual address.
3. The processor recited in claim 2 wherein the mode control mechanism controls an instruction fetch and decode mechanism of the processor so that instructions of said first and second architectures are fetched and aligned for proper decoding by said instruction set management means.
4. The processor recited in claim 3 wherein said memory management unit reads a page mode control bit from a page table entry for said second architecture, said page mode control bit being supplied to said control means to control address translation by said memory management means.
5. The processor recited in claim 1 wherein said control means comprises:
 - an architectural context control mechanism for controlling which architectural context the processor operates under, said architectural context control mechanism controlling an architectural translation method to be used by the memory management unit of the processor when translating an effective address to a virtual address, said architectural context control mechanism further controlling instruction fetch and decode logic of the processor, said architectural context control mechanism further controlling an execution context of process and architected resources, and said architectural context control mechanism further controlling an interrupt and exception mechanism of the processor; and
 - a qualifying mode control mechanism for enabling extensions and limitations to the two architectures, said extensions and limitations allowing a single address translation mechanism to map addresses for one architecture onto translation of addresses of another architecture and a unified interrupt and exception mechanism handling asynchronous interrupts and page translation and protection related exceptions regardless of the architectural context in effect when an interrupt or exception occurs.
6. The processor recited in claim 1 wherein said instruction set management means comprises:
 - first decoding means for decoding instructions for said first instruction set;
 - second decoding means for decoding instructions for said second instruction set;
 - mapping means for mapping decoded instructions from said first decoding means to one or more decoded instructions for said second instructions set; and
 - selection means controlled by said control means for selecting decoded instructions from said mapping means or decoded instructions from said second decoding means.

7. The processor recited in claim 6 wherein said first and second decoding means decode simple instructions of said first and second architectures, simple instructions being instructions that map into a basic operation class and can be handled by a single execution unit, said instruction management means further comprising:

third decoding means for decoding complex instructions of said first instruction set;

fourth decoding means for decoding complex instructions of said second instruction set;

second selection means controlled by said control means for selecting decoded instructions from said third or fourth decoding means;

second mapping means receiving an output of said second selection means for mapping decoded instructions from said third or fourth decoding means for mapping said decoded instructions to multiple, simple instructions of said second instruction set; and

third selection means responsive to a valid signal from one of said third or fourth decoding means for selecting an output of said second mapping means when a complex instruction is decoded or an output of the first mentioned selection means when a simple instruction is decoded.

8. The processor recited in claim 1 wherein said control means comprises a mode control unit which controls the processor to be in one of a plurality of states depending on a state of a plurality of bits stored in registers within said processor including a first bit which determines whether address translation for said second architecture is enabled or disabled, a second bit which indicates whether a current page in memory is for said first or second architectures, and a third bit which is an architecture qualifying bit.

9. The processor recited in claim 8 wherein the processor includes a memory management unit which detects a page mode of instructions read from main memory and informs said mode control unit whether a current page in memory is for said first or second architectures by said second bit.

10. The processor recited in claim 8 wherein the mode control unit further includes a feature control register and is responsive to a power on or reset condition of the processor to initialize said processor in one of first or second modes corresponding to said first and second architectures, respectively, a state

of a bit in said feature control register thereafter enabling dynamic switching between architectural contexts.

11. A computer system comprising

a processor according to claim 1;

an external memory device, said external memory device storing application software for said two architectures; and

a system bus connecting said processor to said external memory device;

said processor having an internal bus connected to said system bus.

12. The computer system recited in claim 11 wherein the processor has a single memory management unit which is implemented using a format common to the two supported architectures.

13. The computer system recited in claim 11 wherein the processor has a single instruction fetch mechanism shared by the two supported architectures and separate instruction decode mechanisms.

14. The computer system recited in claim 11 wherein all execution resources of the processor are common to the two supported architectures.

15. A method implemented in a processor for supporting two separate and distinct instruction-set architectures, said processor operating under a single multitasking operating system, said method comprising the steps of:

controlling which of first or second architectural translation methods is to be used by a memory management unit when translating an effective address to a virtual address in response to a mode control bit supplied by software;

controlling a processor instruction fetch and decode mechanism so that instructions of the two different architectures are decoded properly in response to the mode control bit supplied by software;

translating addresses of the two different architectures by mapping the translation of one architecture onto that of another architecture; and

switching from application software written for one architecture to application software written for another architecture in a multitasking environment.

16. The method implemented in a processor recited in claim 15 further comprising the steps of:

controlling an interrupt and exception mechanism of the microprocessor;

5

enabling extensions and limitations to the two architectures, said extensions and limitations allowing a single address translation mechanism to translate addresses for a first architecture by mapping onto a translation of a second architecture and a unified interrupt and exception mechanism to handle asynchronous interrupts and page translation and protection related exceptions regardless of an architectural context in effect when an interrupt or exception occurred; and

10

15

determining an architectural context under which instructions should execute.

20

25

30

35

40

45

50

55




FIG. 1




FIG.2




FIG. 3




FIG.4




FIG.5

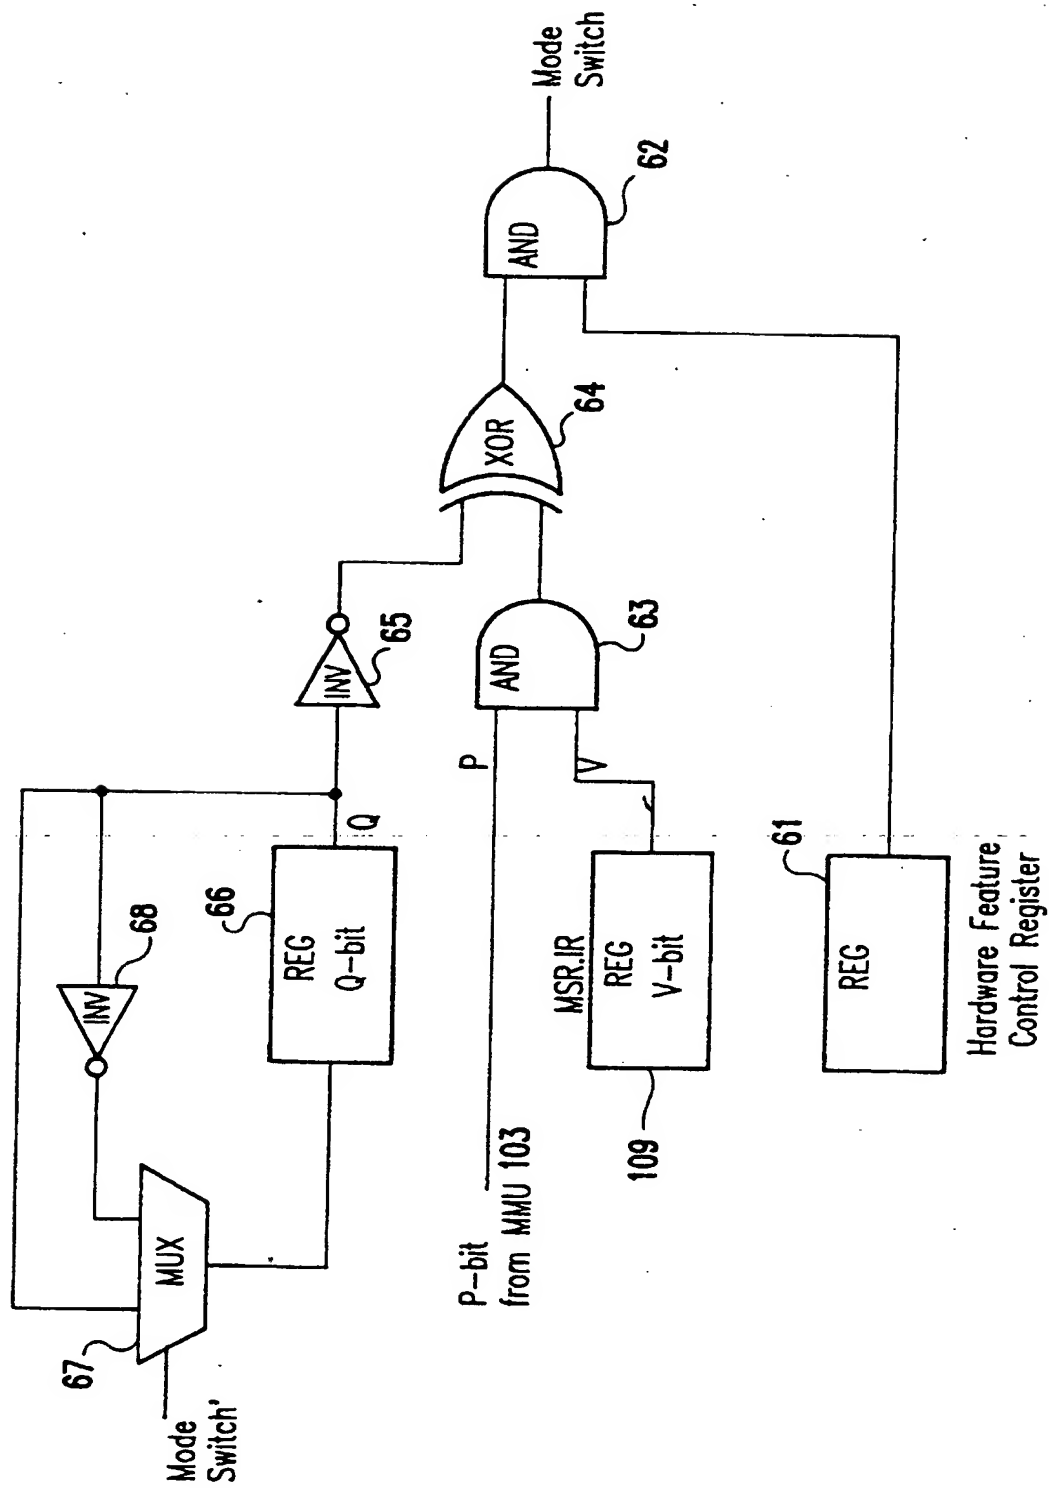


FIG. 6




FIG.7




FIG.8




FIG.9




FIG.10





FIG. 11




1105 $\text{Key} \leftarrow (K_p \ \& \ \text{MSR.PR}) \mid (K_s \ \& \ \sim \text{MSR.PR})$


1106

Key	PP	Access Allowed
0	00	Read/Write
0	01	Read/Write
0	10	Read/Write
0	11	Read Only
1	00	None
1	01	Read Only
1	10	Read/Write
1	11	Read Only

FIG.12

1201  Current Stack Segment Descriptor

1202  CRO

1203  Page Table Entry

1204 ~

WP	CPL	U/S	R/W	Access
0	User	0	0	Read-only
0	User	0	1	R/W
0	User	1	0	None
0	User	1	1	None
0	Super	0	0	R/W
0	Super	0	1	R/W
0	Super	1	0	Read-only
0	Super	1	1	R/W
1	User	0	0	Read-only
1	User	0	1	R/W
1	User	1	0	None
1	User	1	1	None
1	Super	0	0	Read-only
1	Super	0	1	R/W
1	Super	1	0	Read-only
1	Super	1	1	R/W

FIG.13




FIG.14




FIG.15




FIG.16

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 747 808 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
15.01.1997 Bulletin 1997/03

(51) Int Cl.⁶: **G06F 9/318**

(43) Date of publication A2:
11.12.1996 Bulletin 1996/50

(21) Application number: **96480073.4**

(22) Date of filing: **31.05.1996**

(84) Designated Contracting States:
DE FR GB

(30) Priority: **07.06.1995 US 480105**

(71) Applicant: **International Business Machines Corporation**
Armonk, N.Y. 10504 (US)

(72) Inventors:
• **Goetz, John W.**
Jericho, Vermont 05465 (US)

• **Mahin, Stephen W.**
Underhill, Vermont 05489 (US)
• **Bergkvist, John J.**
Williston, Vermont 05495 (US)

(74) Representative: **de Pena, Alain**
Compagnie IBM France
Département de Propriété Intellectuelle
06610 La Gaude (FR)

(54) **Processor capable of supporting two distinct instruction set architectures**

(57) A microprocessor which supports two distinct instruction-set architectures. The microprocessor includes a mode control unit which enables extensions and/or limitations to each of the two architectures and controls the architectural context under which the microprocessor operates. The control unit controls memory management unit (MMU) hardware that is designed to allow address translation to take place under the control of a mode bit so that the translation mechanism can be switched from one architecture to another. A single MMU translates addresses of the two distinct architectures under control of the mode bit which is also used to simultaneously inform instruction decode which architecture is being used so that instructions are properly decoded. The MMU is also capable of mapping the address translation of one architecture onto that of the other so that software written for both architectures may be multi-tasked under the control of a single operating system.




FIG. 8

EP 0 747 808 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 96 48 0073

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	EP-A-0 169 565 (NIPPON ELECTRIC CO) 29 January 1986 * the whole document *	1,6,7, 10,13-15	G06F9/318
A	EP-A-0 272 198 (UNITED TECHNOLOGIES CORP) 22 June 1988 * the whole document *	1,15	
A	EP-A-0 264 215 (AMDAHL CORP) 20 April 1988 * the whole document *	1,2,5, 12,15,16	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
			G06F
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 14 November 1996	Examiner Klocke, L
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 1503 (03.92) (PNC01)